arXiv:1411.3923v1 [cs.CE] 13 Nov 2014

Robust topology optimisation of microstructural details without length scale separation - using a spectral coarse basis preconditioner

Joe Alexandersen*, Boyan S. Lazarov

Department of Mechanical Engineering, Solid Mechanics, Technical University of Denmark, Nils Koppels Allé 404, DK-2800 Kgs. Lyngby, Denmark

23rd of September 2014

Abstract

This paper applies topology optimisation to the design of structures with periodic microstructural details without length scale separation, i.e. considering the complete macroscopic structure and its response, while resolving all microstructural details, as compared to the often used homogenisation approach. The approach takes boundary conditions into account and ensures connected and macroscopically optimised microstructures regardless of the difference in micro- and macroscopic length scales. This results in microstructures tailored for specific applications rather than specific properties.

Dealing with the complete macroscopic structure and its response is computationally challenging as very fine discretisations are needed in order to resolve all microstructural details. Therefore, this article shows the benefits of applying a contrast-independent spectral preconditioner based on the multiscale finite element method (MsFEM) to large structures with fully-resolved microstructural details.

*E-mail: joealex@mek.dtu.dk

The density-based topology optimisation approach combined with a Heaviside projection filter and a stochastic robust formulation is used on various problems, with both periodic and layered microstructures. The presented approach is shown to allow for the topology optimisation of very large problems in MATLAB, specifically a problem with 26 million displacement degrees of freedom in 26 hours using a single computational thread.
Keywords: topology optimisation, multiscale design, microstructure, robust design, multiscale FEM, spectral preconditioner

1 Introduction

The systematic design of novel materials with extremal properties is possible by applying topology optimisation to the design of material microstructures. The usual approach is to consider the homogenised properties of a single unit cell. This completely decouples the problem, as one does not consider the macroscopic response at all. One may then optimise a periodic microstructure to achieve a certain effective property, such as directional stiffness [1, 2], negative Poisson's ratio or thermal expansion coefficient [3], enhanced poroelastic pressure-coupling [4], among many others.

Topology optimisation [5] is an iterative design process which distributes material in a design domain by minimising a selected objective functional, e.g. compliance, material weight or Poisson's ratio, while satisfying a set of constraints. The material distribution is represented by a design field which takes the value one if the point is occupied with material and zero if not. In order to utilise gradient-based optimisation techniques, the design field is allowed to take intermediate values. At each iteration step, the design field is updated using the gradients of the objective and constraint functionals. The discretisation determines the resolution of the optimisation process, as well as its computational cost. Therefore, in order to avoid excessive computations, the design is often limited to a single mechanical element with homogeneous material properties or a single periodic cell.

An alternative multiscale approach to the topological design of material unit cells can be taken by introducing the homogenised macroscopic response in the optimisation process. This is actually the origin of the homogenisation approach to topology optimisation presented in the seminal paper by Bendsøe and Kikuchi [6] and later by Bendsøe [7]. A hierarchical approach is taken in

[8], where the macroscopic density and microstructure is designed iteratively. This approach has been applied to bone modelling [9] and is well suited for parallel computations due to the decoupled homogenisation cells [10]. Similar methodologies have recently been proposed for nonlinear elasticity [11, 12].

The lack of connectivity between the varying microstructural cells is a common problem of the above two-scale optimisation approaches. This may be neglectable when the microstructures are infinitely small, but when considering current manufacturability, a finite size must be attributed to the microstructures. Ensuring connected microstructures has been sparsely treated in the literature, but good examples include [13] in the context of free material optimisation (FMO) and [14] in the context of functionally graded materials (FGM).

The convergence of periodic structures towards homogenised microstructures is investigated in [15, 16]. These papers show that an optimised periodic structure, with finite geometric periodicity, converges to optimised material unit cells with an infinite geometric periodicity, obtained using homogenisation for a single microstructure repeated throughout the computational domain.

This paper restricts itself to a single-scale optimisation approach, where the periodic microstructural details filling a given domain are optimised for the macroscale response. The complete macroscopic structure and its response is considered, while resolving all microstructural details as compared to the homogenisation approach common to all of the above papers, except [15, 16]. The approach takes boundary conditions into account and ensures connected and macroscopically optimised microstructures regardless of the difference in micro- and macroscopic length scales. Furthermore, the restriction to microstructural details controlled by the size of the coarse mesh and a global volume constraint, implicitly introduces a maximum length scale similar to the results presented in [17]. Such a feature can be utilised for providing manufacturability of the design, e.g. in the structural optimisation of high-rise buildings, where the unit cell models a single building unit (room, office or apartment), or for imposing properties which are not directly related or not included in the physical model due to computational cost, e.g., requiring specified porosity in a scaffold design.

Dealing with the complete macroscopic structure and its response is computationally challenging as very fine discretisations are needed in order to resolve all microstructural details. This work therefore explores the application of a contrast-independent spectral preconditioner based on the multiscale fi-

nite element method (MsFEM), as presented in [18], to large structures with fully-resolved microstructural details. The origins of MsFEM can be traced back to [19, 20] where special multiscale basis functions are utilised for solving elliptic problems. The method has been extended further to be applicable to a wide range of linear and non-linear multiscale problems, e.g. [21, 22]. The main idea is to construct basis functions which provide a good approximation of the solution on a coarse grid [22]. For obtaining good accuracy, the coarse basis functions need to have similar oscillatory behaviour as the fine scale solution, which can be achieved by oversampling techniques [22]. The application of MsFEM to linear elasticity and material jumps isolated in the interior of the coarse elements is reported in [23]. This condition cannot be guaranteed in the topology optimisation process, which limits the applicability of the original MsFEM in topology optimisation. For oscillatory high-contrast material properties, an alternative approach is proposed in [24, 25]. The method constructs several basis functions per coarse node, which are capable of representing the important features of the solution and the accuracy of the approximation is controlled by the dimension of the coarse space.

The layout of the paper is as follows: Section 2 covers the theoretical background of the methods used; Section 3 demonstrates the capabilities of the MsFEM approach through numerical experiments; Section 4 covers the implementation details of the presented approach; Section 5 details numerical experiments with the approach; Section 6 goes into details with several microstructural design examples; and Section 7 covers a discussion and conclusions of the presented work.

2 Theoretical background

2.1 Linear elasticity

The considered problem is static linear elasticity, which is governed by the Navier-Cauchy equation:

$$-\operatorname{div}\boldsymbol{\sigma}\left(\mathbf{u}\right) = \mathbf{f} \tag{1a}$$

$$\boldsymbol{\sigma}\left(\mathbf{u}\right) = \mathbf{C} : \boldsymbol{\varepsilon}\left(\mathbf{u}\right) \tag{1b}$$

where $\boldsymbol{\sigma}$ is the stress tensor, $\boldsymbol{\varepsilon}$ is the linearised strain tensor, $\mathbf{C}$ is the linear elastic stiffness tensor, $\mathbf{u}$ denotes the displacement field and $\mathbf{f}$ is the input to

the system in the form of distributed or concentrated forces. The displacements and the input forces are vector quantities with number of elements equal to the dimensionality of the considered physical space. The considered problems are two-dimensional, however, the presented approach is applicable to three dimensions without any significant amendments. The system occupies a bounded physical domain $\Omega \in \mathbb{R}^2$ with the boundary $\Gamma = \overline{\Gamma_{D_i} \cup \Gamma_{N_i}}$ being decomposed into two disjoint subsets for each component $i = 1, 2$. Γ_{D_i} is the part of the boundary with prescribed zero displacement, $u_i = 0$, and Γ_{N_i} denotes the part with prescribed traction, t_i. The stiffness tensor $\mathbf{C}$ is assumed to be isotropic and has the following form $\mathbf{C}(\mathbf{x}) = E(\mathbf{x})\,\mathbf{C}_0$ where $\mathbf{C}_0$ is the stiffness tensor for a predefined Poisson's ratio, $\nu < 0.5$, and unit Young's modulus, $E_0 = 1$. The Young's modulus $E(\mathbf{x})$ is spatially varying and bounded $E(\mathbf{x}) \in [E_{\min}, E_{\max}]$.

The variational formulation, see e.g. [26], of equation (1) is to find $\mathbf{u} \in V_0$ such that:

$$a(\mathbf{u}, \mathbf{v}) = l(\mathbf{v}) \;\text{ for all } \mathbf{v} \in V_0 \tag{2}$$

where $V_0 = \left\{ \mathbf{v} \in [H^1(\Omega)]^2 \,|\, v_i = 0 \text{ on } \Gamma_{D_i}, i = 1, 2 \right\} \subset V = [H^1(\Omega)]^2$ and the bilinear form a and linear form l are defined as:

$$a(\mathbf{u}, \mathbf{v}) = \int_\Omega (\mathbf{C} : \boldsymbol{\varepsilon}(\mathbf{u})) : \boldsymbol{\varepsilon}(\mathbf{v})\, \mathrm{d}\mathbf{x} \tag{3a}$$

$$l(\mathbf{v}) = \int_\Omega (\mathbf{f} \cdot \mathbf{v})\, \mathrm{d}\mathbf{x} + \int_{\Gamma_N} (\mathbf{t} \cdot \mathbf{v})\, \mathrm{d}s \tag{3b}$$

The weak formulation is discretised using the finite element space $V_h \subset V_0$ with vector-valued shape functions defined on a uniform mesh $\mathcal{T}^h$. The discretization leads to a linear system of equations of the form:

$$\mathbf{K}\mathbf{u} = \mathbf{f} \tag{4}$$

where $\mathbf{K}$ is the so-called stiffness matrix, the vector $\mathbf{u}$ consists of the nodal displacements and the vector $\mathbf{f}$ contains the nodal forces applied to the discrete problem.

2.2 Topology optimisation

Topology optimisation is a material distribution method that seeks to find an optimised structure for a given physical problem with respect to a certain

objective functional subject to design constraints [5]. The physical problem, in this work linear elasticity, is discretised using the finite element method, as described in section 2.1. The design is parametrised by attributing each finite element with a relative density, ρ_e, which determines whether the element is solid, $\rho_e = 1$, or void, $\rho_e = 0$. In order to solve the problem using gradient-based optimisation, the binary density variables are relaxed to continuous variables, $\rho_e \in [0, 1]$.

The Young's modulus of each element is made to be dependent on the element relative density using the modified SIMP approach [5] and thus the element stiffness matrix can be formulated as:

$$\mathbf{K}_e\left(\rho_e\right) = \left(E_{min} + (E_{max} - E_{min}){\rho_e}^p\right)\mathbf{K}_0 \tag{5}$$

where E_{min} is the minimum Young's modulus attributed to void areas in order to avoid ill-conditioning of the stiffness matrix, E_{max} is the maximum Young's modulus attributed to solid areas, ρ_e is the element relative density, p is a penalisation parameter and $\mathbf{K}_0$ is the stiffness matrix for an element with a Young's modulus of 1.

In this work, we consider the minimum compliance problem under a constraint on the volume fraction of solid material. The optimisation problem is posed as follows:

$$\begin{aligned} \underset{\boldsymbol{\rho}\in\mathbb{R}^{n_d}}{\text{minimise:}}\ & f(\boldsymbol{\rho}, \mathbf{u}) = \mathbf{f}^{\mathrm{T}}\mathbf{u} \\ \text{subject to:}\ & g(\boldsymbol{\rho}) = \frac{\boldsymbol{\rho}^{\mathrm{T}}\mathbf{v}}{v_f \mathbf{e}^{\mathrm{T}}\mathbf{v}} - 1 \leq 0 \\ & \boldsymbol{\mathscr{R}}(\boldsymbol{\rho}, \mathbf{u}) = \mathbf{0} \\ & 0 \leq \rho_e \leq 1 \quad \text{for } e = 1, ..., n_d \end{aligned} \tag{6}$$

where f is the compliance functional, g is the volume constraint functional, $\boldsymbol{\rho}$ is a vector of the n_d-number of design variables, $\mathbf{v}$ is a vector of the n_d-number of element volumes, $\boldsymbol{e}$ is a n_d-entry long vector of ones and $\boldsymbol{\mathscr{R}}\left(\boldsymbol{\rho}, \mathbf{u}\right) = \mathbf{K}(\boldsymbol{\rho})\,\mathbf{u} - \mathbf{f}$ is the residual of the discretised system of equations, . The optimisation problem is solved using the nested formulation, where the discretised system of equations for the state field is solved separately from the design problem.

2.3 Computational issues of topology optimisation

Topology optimisation is an iterative process and requires the solution of a large system of equations for the state field, equation (4), at every design iteration. Solving the linear systems of equations can often account for more than 99% of the total computational time [27]. Therefore, the solution cost of the optimisation procedure depends strongly on the time necessary for solving the state equations. Realistic mechanical problems with microstructural scales comparable to the macroscale require very fine discretisations to capture the microstructural scales, making direct solvers prohibitive even in two dimensions. An alternative is to utilise iterative solvers, where the convergence and the solution cost is controlled mainly by the applied preconditioner. Classical preconditioning techniques, such as incomplete factorisation, sparse approximate inverse or stationary iterative methods, cannot provide a mesh-independent number of iterations. Furthermore, for topology optimisation problems with high contrast between the material parameters, E_{max}/E_{min}, the number of iterations increases with increasing contrast [27, 28]. Geometric multigrid techniques [29] can provide mesh-independent solution for smooth material properties, however, similar to the classical preconditioners, the number of iterations depends on the contrast when the mesh is not aligned with the different materials.

2.4 Multiscale finite element (MsFEM) coarse basis for linear elasticity

An effective alternative to geometric multigrid and classical preconditioners was demonstrated in [24, 25]. The method is a multiscale approach providing the solution of the diffusion equation with strongly heterogeneous coefficients. The solution cost is independent of the contrast of the diffusion coefficients and is significantly reduced by using coarse space approximations which contain important features of the solution. The approach has been extended and applied in topology optimisation of linear elastic problems in [18] and the main steps are outlined below in the context of structures with periodic microstructural details.

The fine mesh $\mathcal{T}^h$, utilised for discretising equation (1) in section 2.1, is assumed to be obtained by a refinement of a coarser one $\mathcal{T}^H = \{K_j\}_{j=1}^{N_c}$, where K denotes a coarse mesh cell and N_c is the number of coarse nodes. In the context of microstructural design, the coarse cells represent the unit

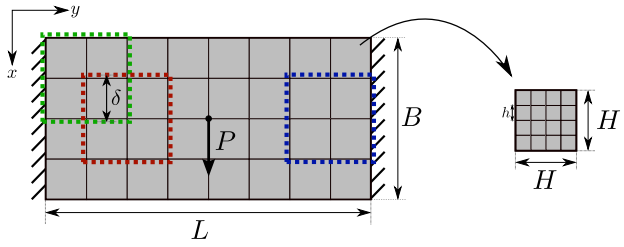

Figure 1: Illustration and dimensions of a double-clamped beam subjected to a concentrated load in the centre. The beam is made up of a single periodic microstructure with square unit cells. The fine mesh, coarse mesh and agglomerates are illustrated.

cells and the fine cells the discretisation of the unit cells, which is illustrated in figure 1.

The nodes of the coarse mesh are denoted as $\{\mathbf{y}_i\}_{i=1}^{N_c}$ and the neighbourhood of node $\mathbf{y}_i$ is defined as:

$$\omega_i = \bigcup \{\overline{K_j} \in \mathcal{T}^H; \mathbf{y}_i \in \overline{K_j}\} \tag{7}$$

which is illustrated in figure 1 using dotted lines. These neighbourhoods will be denoted as agglomerates, since they can be viewed as a group of coarse elements agglomerated together with an overlap δ [30].

A set of coarse basis functions, $\{\boldsymbol{\phi}_{i,j}, j = 1 \dots N_{i,j}\}$, defined with respect to $\mathcal{T}^h$, is introduced for each node $\mathbf{y}_i$ in the coarse mesh with support on ω_i, where $N_{i,j}$ is the number of basis functions for the i'th coarse node. The coarse basis functions are based on the modes obtained by solving the following local generalised eigenvalue problem on each agglomerate ω_i:

$$-\operatorname{div}(\mathbf{C}(\mathbf{x}) : \boldsymbol{\varepsilon}(\mathbf{u})) = \lambda \mathrm{E}(\mathbf{x})\mathbf{u}, \quad \mathbf{x} \in \omega_{\mathrm{i}} \tag{8}$$

The local eigenvalue problem is discretised using the fine mesh, which leads to the following discrete generalised eigenvalue problem in matrix-vector form:

$$\mathbf{K}_{\omega_i}\boldsymbol{\psi}_j^{\omega_i} = \lambda_j^{\omega_i}\mathbf{M}_{\omega_i}\boldsymbol{\psi}_j^{\omega_i} \tag{9}$$

where $\mathbf{K}_{\omega_i}$ is the stiffness matrix, $\mathbf{M}_{\omega_i}$ is a stiffness-based mass matrix, $\boldsymbol{\psi}_j^{\omega_i}$ is the j'th eigenvector and $\lambda_j^{\omega_i}$ is the j'th eigenvalue, all for the i'th agglomerate, ω_i. It is important to note that the local matrices take the physical

boundary conditions of the full problem into account and the coarse basis thus automatically fulfils the specified boundary conditions.

The eigenvalues are ordered as $\lambda_1^{\omega_i} \leq \lambda_2^{\omega_i} \leq \cdots \leq \lambda_j^{\omega_i} \leq \ldots$ and the first several eigenvectors corresponding to eigenvalues smaller than a prescribed global threshold, λ_Ω, are selected to form the coarse basis. The coarse basis functions, represented on the fine grid, are defined as $\boldsymbol{\phi}_{i,j} = \boldsymbol{\psi}_j^{\omega_i} \chi_i$; that is they are constructed by multiplying the eigenfunctions, $\boldsymbol{\psi}_j^{\omega_i}$, with a partition of unity, $\{\chi_i\}_{i=1}^{N_c}$, subordinated to ω_i such that $\chi_i \in H^1(\Omega)$ and $|\nabla \chi_i| \leq 1/H, i = 1, \ldots, N_c$, and H is the characteristic length of a coarse element K. Therefore, the set of coarse basis functions $\left\{\boldsymbol{\phi}_{i,j}\right\}$ associated with node $\mathbf{y}_i$ is defined as the fine space finite element interpolant of $\left\{\boldsymbol{\psi}_j^{\omega_i}(\mathbf{x})\, \chi_i(\mathbf{x}), j = 1 \ldots N_i\right\}$, see e.g. [31]. For each ω_i, N_i is determined as the number of eigenvalues smaller than a globally selected threshold value, λ_Ω.

The solution in the coarse space is sought as $\mathbf{u}_a = \sum_{i,j} c_{i,j} \boldsymbol{\phi}_{i,j}$, where $c_{i,j}$ is the coefficient of the j'th coarse basis function for the i'th agglomerate, ω_i. A coarse discretisation of the variational formulation is given as $\mathbf{K}_c \mathbf{u}_c = \mathbf{f}_c$, where the coarse stiffness matrix $\mathbf{K}_c$ and the coarse right-hand side $\mathbf{f}_c$ are obtained as:

$$\mathbf{K}_c = \mathbf{R}_c \mathbf{K} {\mathbf{R}_c}^{\mathrm{T}}, \quad \mathbf{f}_c = \mathbf{R}_c \mathbf{f} \tag{10}$$

The vector $\mathbf{u}_c$ contains the coefficients $c_{i,j}$ and ${\mathbf{R}_c}^{\mathrm{T}} = \left[\boldsymbol{\phi}_1, \boldsymbol{\phi}_2, \ldots, \boldsymbol{\phi}_{N_t}\right]$, where $N_t = \sum_{i=1}^{N_c} N_i$, is a matrix describing the mapping from the coarse to the fine space and consisting of nodal values of the coarse basis functions in the fine space. An approximation of the nodal solution in the fine space is obtained as $\mathbf{u}_a = {\mathbf{R}_c}^{\mathrm{T}} \mathbf{u}_c$.

2.4.1 Periodic microstructure

Building the MsFEM basis for general linear elastic systems requires solving the eigenproblem, equation (9), for each agglomerate and as discussed in [18] this process is computationally very expensive. It is important to note that when dealing with periodic microstructural details, the computation of eigenfunctions is limited to a small amount of unique agglomerates. Figure 1 shows the three unique agglomerates for a double-clamped beam with a periodic microstructure. Due to the imposed periodicity of the microstructure, all agglomerates along the left-hand side are the same (constrained in both x- and y-direction at the left-hand side), all agglomerates along the right-hand

side are the same (constrained in both x- and y-direction at the right-hand side) and all internal agglomerates are the same (unconstrained and allowing rigid body motion). This means that significant time can be saved on the computation of the local spectral bases.

2.4.2 Application as multigrid-like preconditioner

MsFEM can be applied as a solver in topology optimisation, however, if the coarse solver does not provide an accurate enough approximation of the fine-scale solution, the optimisation will results in sub-optimal designs with disconnected features [18]. An alternative is proposed in [24], where the basis is utilised in a two-level additive Schwarz preconditioner for the iterative solution of the fine scale solver. The preconditioner is scalable and results in contrast independent number of iterations. Here, the coarse basis is simply applied as a coarse grid solve in a multigrid-like preconditioner [30] for the iterative solution of the fine-scale system of equations, equation (4). The residual of the fine-scale system is restricted to the coarse basis and the correction is approximated using a MsFEM coarse-scale solve and projected back to the fine space. In the current work, the MsFEM coarse-scale solve is combined with a single pre- and post-smoothing using symmetric Gauss-Seidel and used as a preconditioner for the generalised minimal residual (GMRES) iterative method [32]. The linear elasticity problem, as well as the preconditioner, is symmetric and positive definite and, hence, the preconditioned conjugate gradient (PCG) method could be used. However, the authors' experience is that GMRES performs better compared to PCG for the presented problems and thus GMRES is used throughout.

3 Demonstration of the MsFEM preconditioner

3.1 Numerical experiments

In order to demonstrate the performance and applicability of MsFEM for topology optimisation problems, a test problem is formulated. The test problem is a short beam clamped at both ends and subjected to a concentrated load in the centre, as shown in figure 1. The height of the beam is set to $B = 1$, the length to $L = 2$ and the size of the force is set to $P = 0.01$. The beam is made up of 8 by 16 square coarse cells, with an edge length of $H = 1/8$, containing the same microstructure. Each coarse cell is discretised

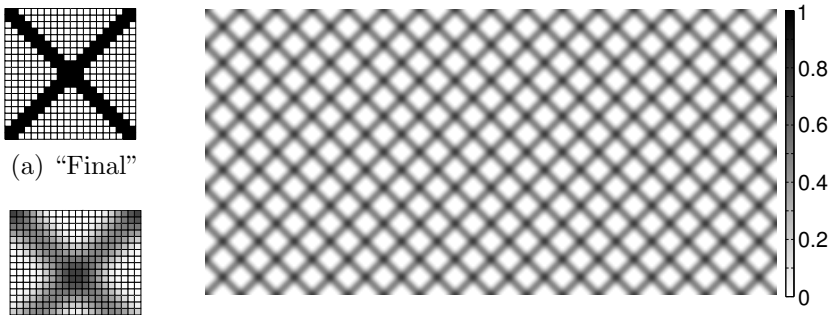

(a) "Final"

(b) "Intermediate"

(c) Macrostructure: "Intermediate"

Figure 2: The investigated cross structure with a unit cell discretised by 20x20 elements. The macrostructure is made up of 8 by 16 unit cells, which is shown in two versions: (a) final 0-1 design and (b) the corresponding "intermediate" design generated by filtering the first. Subfigure (c) shows the macrostructure for the "intermediate" design.

with 20 by 20 linear finite elements, leading to a full macroscopic mesh of 160 by 320. The problem is investigated for varying values of the eigenvalue threshold, λ_Ω, and MsFEM is applied both as a solver and as a preconditioner in combination with GMRES.

Figure 2 shows the structure under consideration for these numerical experiments. The microstructure is chosen as a cross structure and is considered in two versions, the first of which is the clear 0-1 design shown in figure 2a. This design is representative of a final design of a topology optimisation process using a Heaviside projection as will be described in section 4.3. The structure is also considered in a smeared out version, as can be seen in figure 2b, which is representative of an intermediate design during the optimisation process where the design usually contains significant amounts of intermediate stiffnesses. The smeared out design is obtained using a density filter, as will be described in section 4.4, with a filter radius of 4 times the element side length. The macrostructure for the intermediate design case is shown in figure 2c. The structure is analysed with a SIMP penalisation parameter of p=3. Similar trends are observed for the two structures and thus only the results for the filtered version are shown.

Figure 3 shows the relative error in the energy norm, $\|\mathbf{u} - \mathbf{u}_a\|_{\mathbf{K}} / \|\mathbf{u}\|_{\mathbf{K}} =$

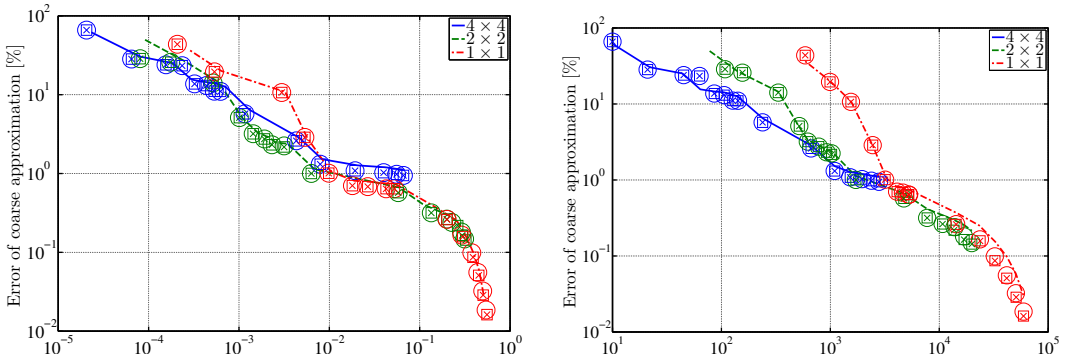

(a) Eigenvalue threshold, λ_Ω (b) Number of coarse DOF

Figure 3: Relative error in the energy norm for a single MsFEM solve of the filtered cross structure. The error is shown as a function of (a) varying eigenvalue threshold and (b) the corresponding number of coarse degrees of freedom. Colours denote different agglomerate sizes, see legend, and symbols denote different contrasts: lines; $E_{min} = 10^{-3}$, circles; $E_{min} = 10^{-6}$, squares; $E_{min} = 10^{-9}$ and crosses; $E_{min} = 10^{-12}$.

$(\mathbf{u} - \mathbf{u}_a)^{\mathrm{T}}\mathbf{K}(\mathbf{u} - \mathbf{u}_a)/\mathbf{u}^{\mathrm{T}}\mathbf{K}\mathbf{u}$, of the solution obtained from a single coarse-scale MsFEM solve for the intermediate design case. The analysis is done for a constant maximum Young's modulus, $E_{max} = 1$ and a set of minimum Young's moduli, $E_{min} \in \{10^{-3}, 10^{-6}, 10^{-9}, 10^{-12}\}$, as well as for different sizes of the agglomerates, denoted by the numbers of unit cells covered by each coarse cell, $\{1 \times 1, 2 \times 2, 4 \times 4\}$. It can be seen that the approximated solution converges with respect to both the eigenvalue threshold and the corresponding number of degrees of freedom. Even though small differences are observed between $E_{min} = 10^{-3}$ and the lower, it can clearly be seen that the behaviour is contrast-independent for $E_{min} \in \{10^{-6}, 10^{-9}, 10^{-12}\}$. It can be seen that the energy norm error follows a common trend, in relation to the eigenvalue threshold, independent of the agglomerate size, similar to the scalar case [25]. Clear differences can be seen in the lower end of the datasets for the error in relation to the number of coarse degrees of freedom. The curves do, however, converge towards a common trend as the number of degrees of freedom is increased.

As can be seen from the error plots, a very large set of coarse basis functions need to be used to obtain a good accuracy of a single MsFEM solve. This results in a very large computational cost, both for solving the

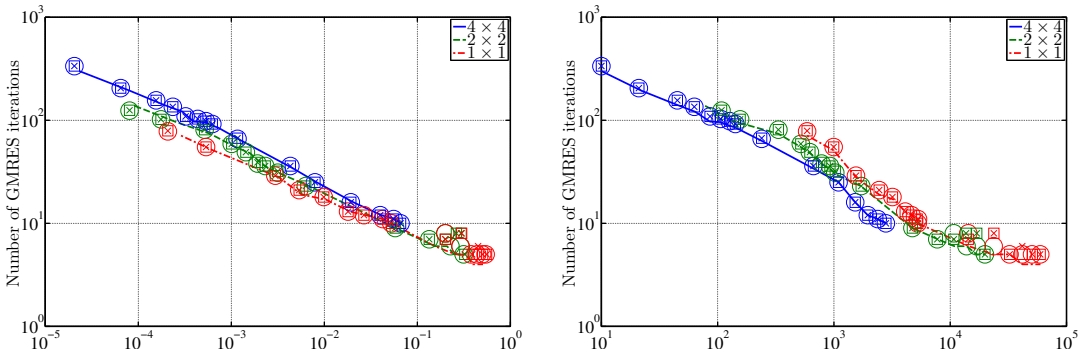

(a) Eigenvalue threshold, λ_Ω (b) Number of coarse DOF

Figure 4: Number of GMRES iterations until convergence for the filtered cross structure. The number of iterations is shown as a function of (a) varying eigenvalue threshold and (b) the corresponding number of coarse degrees of freedom. See figure 3 for explanation of symbols and colours.

generalised eigenvalue problems but also for solving the coarse-scale system. This is why the MsFEM basis is applied as a preconditioner to the full system instead.

Figure 4 shows the number of GMRES iterations needed to converge to a precision of $\varepsilon_{rel} = 10^{-8}$ relative to the preconditioned norm of the forcing vector. As expected, the number of GMRES iterations decreases as the eigenvalue threshold is increased and the basis is enlarged. Besides the results showing clear contrast-independence, they also show that the convergence of the GMRES solver is not significantly affected by the size of the agglomerate. More specifically, by looking at figure 4a it can be seen that by choosing the same eigenvalue threshold, the number of GMRES iterations needed to converge is very close to the same.

3.2 Diagonal weighting matrix

Instead of computing the stiffness-based mass matrix for a given agglomerate, as specified by equation (8), it is possible to simply use the diagonal of the agglomerate stiffness matrix as the weighting matrix of the generalised eigenproblem, equation (9). This is because the two weighting matrices are spectrally equivalent and $\lambda_{\text{diag}(\mathbf{K})} \approx Ch^2\lambda_{\mathbf{M}}$ for some constant C [29, 33]. This is interesting from a computational point of view, as the diagonal leads

to a cheaper solution procedure; both in that the stiffness-based mass matrix is not assembled, and in that the matrix-products involved in solving the eigenproblems become cheaper. The numerically observed convergence behaviour is very similar to that seen in figure 3 and is thus not shown. The diagonal of the agglomerate stiffness matrix is used as the weighting matrix throughout the rest of this paper.

3.3 Spectral basis computational cost

Considering the total computational time taken for the projection operation, equation (10), and the iterative solution, there exists an optimal basis size that provides the fastest solution process. Many factors influence this time, such as the sparsity of the projection matrix and also the number of iterations taken to reach convergence. It should also be noted that the times are not the same for the three agglomerate sizes, because the larger the agglomerate, the larger the support area and the denser the projection matrix becomes. Thus, it is clear that increasing the size of the basis does not necessarily lead to a faster solution even though it results in fewer GMRES iterations, as each of these become more expensive along with the projection operation. It can thus be concluded that it cannot pay to increase the size of the agglomerate, as it does not appear to influence the error nor the convergence of the iterative solver, figure 4, as well as leading to a more expensive solution procedure due to larger eigenproblems and denser projection matrices.

3.4 Note on length scale separation

The main assumption in standard homogenisation is that the macroscale medium consists of an infinite number of periodic unit cells. The homogenised solution is an asymptotic solution, which is valid only when the cell characteristic length H is orders of magnitude smaller than the macroscopic problem scale, i.e., $H << L$. For cell characteristic lengths in the intermediate range, $H < L$, the homogenised solution can differ significantly from the true solution. Also, the homogenised material properties cannot account for the boundary conditions and any load variations of the fine scale problem.

Furthermore, in the design of functionally graded materials, i.e. materials with spatially varying properties, the periodicity condition is often violated. The transition from one cell type to another cell type is applied abruptly within a distance comparable to H. In the above mentioned cases,

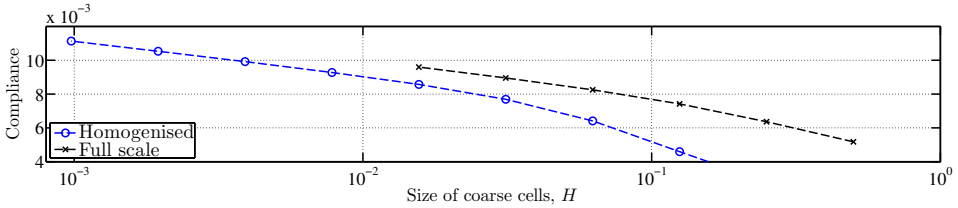

(a) Double-clamped beam with concentrated load

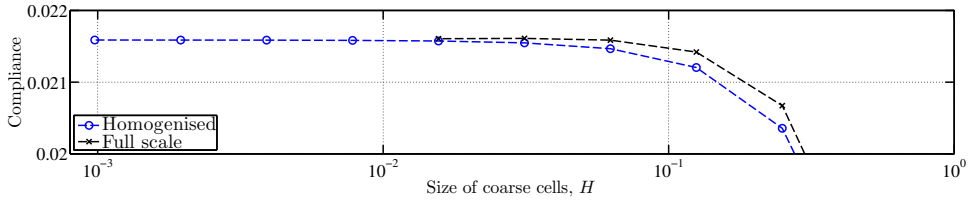

(b) Cantilever beam with distributed load

Figure 5: Convergence of compliance with respect to the size of the coarse cell using both homogenisation and full scale analysis using MsFEM-GMRES.

the proposed MsFEM-GMRES approach will provide the physical response without any simplifications for relatively low computational cost. Hence, it will be superior, but more costly, compared to the standard homogenisation for optimisation problems, where localised constraints (loading, boundary conditions, surface effects) significantly affect the design topology.

Figure 5 shows the change of compliance with respect to the size of the coarse cell using both homogenisation and full-scale analysis using MsFEM-GMRES. The compliance is shown both for the double-clamped beam with a concentrated load in the centre, shown in figure 1, as well as a cantilever with a distributed load at the end, as will be introduced in section 6.2. It should be noted that convergence is not expected for the case with a point-load, which is also observed for both methods in figure 5a. It is important to note that any observed convergence is slightly different for the two solution types. For the full-scale analysis, the size of the microstructure decreases when decreasing the size of the coarse cells. The convergence of compliance with respect to microstructure size is assumed to be dominant over the increased finite element resolution of the macroscopic domain. For the homogenised case, decreasing the size of the coarse cells solely increases the resolution of the finite element discretisation of the homogenised problem. That is, the microstructure size can be seen as constant and is assumed to be small enough

to fulfil the assumptions of homogenisation. Thus, figure 5 illustrates that the macroscopic compliance of structures with finite size periodic microstructural details converges towards that estimated by homogenisation. However, it also emphasises that if one attributes a large finite size to a microstructure designed using homogenisation, one may be introducing a substantial error in the macroscopic response. This is especially pronounced for problems with localised behaviour as the illustrated double-clamped beam with a concentrated load in the centre, figure 5a.

4 Implementation details

4.1 Topology optimisation of microstructural details

Referring to figure 1, the design domain is made up of repeated unit cells covering the entire domain. The unit cells contain a periodic microstructure that is discretised with a certain number of finite elements and each of these are assigned a relative density which is controlled by a corresponding design variable. Due to the imposed periodicity of the design, the same design variable controls the relative density of several fine-scale elements in the macroscopic problem. The total sensitivity for a given design variable is thus calculated by summing up local sensitivities of all fine-scale elements which it controls.

It is important to emphasise that the state field is solved on the macroscopic level resolving all microstructural details. The objective functional is thus the fine-scale macroscopic compliance. By considering the full structure, one ensures that all microstructural details will be connected in order to provide a beneficial macroscopic behaviour, as disconnected members would be detrimental to the compliance objective.

One of the aims of this paper is to investigate the performance of the spectral basis preconditioner, as described in section 2.4, and the performance is thus compared to reference results. For calculating reference results, the equation system is solved using the standard backslash (`mldivide`) in MATLAB and these are then compared to the results obtained when using the spectral preconditioner based on MsFEM combined with the standard MATLAB implementation of the GMRES iterative method (`gmres`).

The optimisation problem is solved using the method of moving asymptotes (MMA) [34], of which the used MATLAB implementation is courtesy of

Krister Svanberg.

4.2 Adjoint sensitivity analysis

In order to use a gradient-based optimisation method, such as MMA, one needs the sensitivities, or gradients, of the objective functional and any constraint functionals. These sensitivities are easily found for the structural compliance using the discrete adjoint method [35]. Because the problem is self-adjoint, the sensitivity for the e'th element simply becomes:

$$\frac{\partial f}{\partial \rho_e} = -\mathbf{u}_e \frac{\partial \mathbf{K}_e}{\partial \rho_e} \mathbf{u}_e \tag{11}$$

where $\mathbf{u}_e$ is the element nodal displacements and $\mathbf{K}_e$ is the element stiffness matrix. The derivative of the element stiffness matrix with respect to the element density is easily found as:

$$\frac{\partial \mathbf{K}_e}{\partial \rho_e} = p(E_{max} - E_{min})\rho_e^{p-1}\mathbf{K}_0 \tag{12}$$

4.3 Heaviside projection filtering

Heaviside projection filtering [36, 37] is applied in order to facilitate the formation of crisply-defined topologies, or 0-1 topologies, without intermediate densities. First, the design variables are filtered using density filtering [38] and subsequently the filtered density variables are passed through a smoothed Heaviside function in order to project the densities below or above a given threshold to 0 or 1, respectively:

$$\bar{\tilde{\rho}}_e = \frac{\tanh(\beta\eta) + \tanh(\beta(\tilde{\rho}_e - \eta))}{\tanh(\beta\eta) + \tanh(\beta(1-\eta))} \tag{13}$$

where $\eta \in [0;1]$ determines the projection threshold and $\beta > 0$ determines the sharpness of the approximation.

After the design variables, $\boldsymbol{\rho}$, have been filtered, $\tilde{\boldsymbol{\rho}}$, and projected, $\bar{\tilde{\boldsymbol{\rho}}}$, one has obtained the actual physical densities. These physical densities are used for the calculation of the element stiffness matrix, equation (5), and therefore, the sensitivities, equation (11), are corrected using the chain rule, as detailed in e.g. [39].

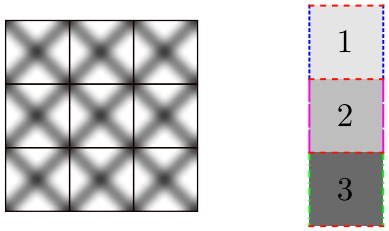

(a) Periodic filtering (b) Multiple blocks

Figure 6: Illustration showing the filtering procedures: (a) single periodic microstructure where the filtering is performed on a 3×3 block - (b) multiple blocks where periodicity and connectivity constraints imposed on the filtered density field. The colours and line styles indicate which edges are coupled together explicitly by collapsing their degrees of freedom.

The standard neighbourhood-based filtering approach, see e.g. [39], is used for all designs in this paper, except for the problems containing multiple microstructures, as will be described in section 4.5. For these problems, the PDE-based density filter, as introduced in [40], is used due to the fact that connectivity and periodicity boundary conditions are easily introduced. For the single microstructure designs, periodicity is introduced by applying the standard filter on a 3×3 block of the microstructural design as illustrated by figure 6a.

4.4 Robust topology optimisation

All manufacturing methods introduce sources of uncertainty in the final realised structure, the most prevalent of which is an uncertainty in the manufactured geometry. This is caused by either too little or too much material being removed or added during the manufacturing process. The possible variability in the manufactured structures can be taken into account by considering several design realisations during the optimisation process [37].

The optimisation problem is cast in a stochastic robust formulation [41], in which the mean and variance of the macroscopic compliance is considered. The different realisations of the design are obtained by using several projections using different threshold values simulating uniform over- and under-

material deposition or removal. The optimisation problem is posed as follows:

$$\begin{aligned} \underset{\boldsymbol{\rho}\in\mathbb{R}^{n_d}}{\text{minimise:}}\ & f(\boldsymbol{\rho},\mathbf{u}) = \mathrm{E}\left[\mathbf{f}^{\mathrm{T}}\mathbf{u}\right] + \kappa\sqrt{\mathrm{Var}\left[\mathbf{f}^{\mathrm{T}}\mathbf{u}\right]} \\ \text{subject to:}\ & g(\boldsymbol{\rho}) = \frac{\mathrm{E}\left[\tilde{\bar{\boldsymbol{\rho}}}^{\mathrm{T}}\mathbf{v}\right]}{v_f\mathbf{e}^{\mathrm{T}}\mathbf{v}} - 1 \leq 0 \\ & \mathscr{R}(\boldsymbol{\rho},\mathbf{u}) = \mathbf{0} \\ & 0 \leq \rho_i \leq 1 \quad \text{for } i = 1,...,n_d \end{aligned} \tag{14}$$

where $\mathrm{E}[\cdot]$ and $\mathrm{Var}[\cdot]$ denotes the expected value and variance of a given quantity, respectively, and κ is a weighting factor for the inclusion of the variance in the objective. The sensitivities are computed as described in [41].

Throughout this paper, the weighting factor, κ, is set to 1.0 and three design realisations are used with equal weights, specifically the set $\eta \in \{0.3, 0.5, 0.7\}$. The final designs are verified using a large set of uniformly distributed points, $\eta \in [0.3, 0.7]$.

4.4.1 Constant β approach for moderate β

A constant β approach as presented in [42] is adopted in order to eliminate changes in the allowable minimum length scale during the optimisation process. By having a constant β, controlling the slope of the projection, in combination with a constant filter radius and constant thresholds, one ensures a constant minimum length scale. This ensures that small features that cannot be supported by the mesh and length scale do not appear during the optimisation process. Using the classical β-continuation, these smaller features are progressively removed during the optimisation process and convergence issues can be encountered along the way. The approach relies on simple changes to the size of the initial asymptotes in the MMA algorithm, which is set to $\frac{0.5}{1+\beta}$. No external move limits are imposed and thus the full control of the updates are left in the hands of MMA. Even though one removes the need to use a continuation approach on the β-parameter by making these modifications, it is observed to be beneficial to perform continuation of the SIMP penalisation parameter in order to get high quality solutions, as also noted in the original paper [42]. If starting from a homogeneous intermediate guess, the constant beta approach retains the possibility to relatively freely form topologies. However, if starting from a random guess, many areas will

be projected to solid and void and the constant beta approach then in many ways resembles some form of level set approach.

For the single design results presented in this work, a continuation approach is used on the SIMP penalisation parameter, increasing it in steps of $p \in \{1, 2, 3, 4, 5\}$ when the maximum relative change in objective function reaches below 0.001 or when the number of iterations for the same parameter size reaches 50. After the last step, the optimisation is stopped when a maximum number of iterations is reached or the maximum relative change criteria is met.

Unfortunately, when combined with the robust topology optimisation formulation, one is somewhat restricted in the maximum size of β. The size of β needs to be moderate in order to ensure the ability of the optimiser to form a good initial topology from both homogeneous and random starting distributions. The important thing is for the smooth parts of the shifted approximative Heaviside functions to overlap, or rather that the parts with non-negligible sensitivities overlap, and this sets an upper limit for the initial β for the optimisation to proceed in practice. Thus, the robust topology optimised designs in this paper have been obtained from a three step continuation approach. During the initial stage where the topology is forming, the SIMP penalisation parameter is set to $p_0 = 1.5$ and β is set to a moderate value of $\beta_0 \in [8, 32]$ depending on the filter radius - the larger the filter radius relative to the element size, the larger β is needed to provide the same level of intermediate densities [36]. During the next stage, the SIMP penalisation parameter is raised to $p_1 \in [3.0, 5.0]$ depending on the problem - some microstructural design problems require higher penalisation than others to form 0-1 topologies. For the final stage of the optimisation, the β parameter is increased to $\beta_1 \in [32, 128]$, again depending on the filter radius. Even though one has not completely eliminated the continuation approaches, the fact that β is already set to a moderately high value from the beginning reduces the likelihood of unsupportable features forming during the optimisation process. The next step is introduced after 100 iterations or when the maximum relative change in objective function reaches below 0.001. After the last step, the optimisation is stopped when a maximum number of iterations is reached or the maximum relative change criteria is met.

4.5 Multiple blocks

Restricting the design to be made up of a single periodic microstructure is quite a severe constraint on the design freedom. The obtained designs performance will be far from the unconstrained optimal design. This is due to the fact that the same microstructure is placed throughout all parts of the design domain and the optimised microstructure will thus be one that satisfies optimality in an average sense, but where the parts with the largest sensitivities will dominate the optimisation procedure. This restriction can be partially alleviated by partitioning the design domain into multiple subdomains, within which a unique locally periodic microstructure is designed.

Smoothly connected microstructures, both within the subdomains and across their interfaces, are obtained by considering the full macroscopic problem and by imposing periodicity and connectivity constraints on the optimised design. This is achieved by using the PDE-based density filter [40], which easily allows for imposing such constraints on the filtered density field by collapsing the corresponding degrees of freedom in the discrete filter problem. Figure 6b illustrates the periodicity and connectivity constraints imposed on the filtered density field for a design with three layers of locally periodic microstructures. The colours indicate which edges are coupled together explicitly by collapsing their degrees of freedom. The top and bottom (red) of all cells are to be connected ensuring both periodicity, in the vertical direction, within the block and connectivity across the interfaces. The left and right sides (blue, green and magenta) of each cell type are to be connected ensuring periodicity within the blocks in the horisontal direction.

The increase in the number of microstructures present in the design, expands the number of unique agglomerates to consider when calculating the eigenmodes for the spectral basis. For instance, considering a double-clamped beam where the design is made up of three layers with each their microstructure, the number of agglomerates increases to 15. There is 3 sets to accommodate the different boundary conditions, as discussed for the single microstructure case in section 2.4.1, within which there now exists 5 unique agglomerate types, more specifically one for each layer and one for each interface.

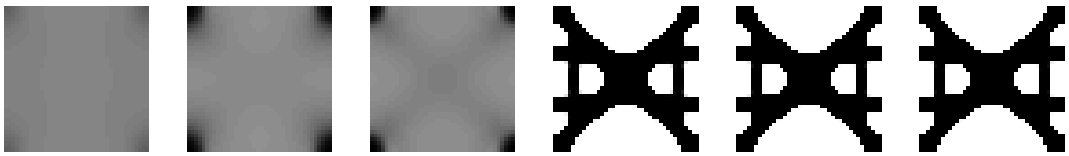

(a) Iteration 1 (b) Iteration 10 (c) Iteration 30 (d) Iteration 170 (e) Iteration 190 (f) Iteration 200

Figure 7: The design distributions for the first and final stages of the optimisation process; design iterations 1, 10, 30, 170, 190 and 200.

5 Numerical experiments

5.1 Basis re-utilisation

It is rather trivial to argue that while the design is evolving slowly, which is prevalent during the first and final stages of the presented optimisation procedure, it is possible to reuse the previously computed spectral basis for the iterative solution. In order to test this hypothesis, the previously described double-clamped beam problem, as shown in figure 1, is optimised using different strategies for updating the MsFEM basis during the optimisation process. The beam consists of 4 by 8 coarse cells within which the microstructure is discretised using 40 by 40 elements. The investigation is carried out for a constant $\beta = 64$ using a single realisation with threshold $\eta = 0.5$. The maximum number of design iterations is set to 200 in order to keep the total number of iterations the same for all the different strategies in order to conduct a fair comparison.

Figure 7 illustrates the relatively slow development of the design during the first and final stages of the optimisation. It can be seen that the design is evolving quite slowly during the first iterations, figures 7a to 7c, due to the low SIMP penalisation value and this supports the arguments that the MsFEM basis does not need to be updated very frequently during this period. Furthermore, it is clearly shown that the design is hardly changing during the final iterations, figures 7d to 7f, where the topology has been determined and the final minor adjustments are being carried out.

5.2 Simple update scheme based on GMRES-iterations

In [18] the spectral basis is updated for the agglomerates within which the design is changing more than a specified tolerance. Here a simple heuristic update scheme based on the number of GMRES iterations is used. The spectral basis is updated when the number of GMRES iterations changes more than a given percentage from the design iteration where the basis was previously computed.

In order to further speed up the iterative solution of the linear systems of equations, the GMRES solver can be initialised with the displacement solution from the previous design iteration as the initial guess for the GMRES solver for the new design iteration, rather than starting from a zero-vector initial guess. When initialisation of GMRES is used, the heuristic basis update rule is modified so that if the number of GMRES iterations decreases after a basis update, the update threshold is updated to be with respect to the lowest number of GMRES iterations encountered after the previous basis computation.

Furthermore, it has been shown in [43, 28] that the convergence criteria can be relaxed significantly during topology optimisation when using an iterative solver. The reason for this is that one does not necessarily require the displacement solution to be solved to full precision in order for the optimisation to progress correctly. This leads to significant reductions in computation time compared to solving to full precision. For the investigations presented in this section, the stopping tolerance is set to $\varepsilon_{rel} = 10^{-5}$ relative to the preconditioned norm of the forcing vector.

A metric that is correlated to the development of the design is the non-discreteness measure, M_{nd}, as introduced in [39]:

$$M_{nd} = \frac{\sum_{i=1}^{n_d} 4\bar{\tilde{\rho}}_i(1 - \bar{\tilde{\rho}}_i)}{n_d} \tag{15}$$

M_{nd} is an indication of the amount of intermediate densities and thus an indication of the contrast of the stiffness distribution. When M_{nd} is high, there is large amount of intermediate densities and when M_{nd} is low, there is small amount of intermediate densities. For volume-constrained compliance minimisation where the design will never contain all 0 or all 1 densities, this is equivalent to saying: when M_{nd} is high, there is low contrast and when M_{nd} is low, there is high contrast. The top plot of figure 8 shows how M_{nd} changes during the optimisation process. It is clearly seen that the level of

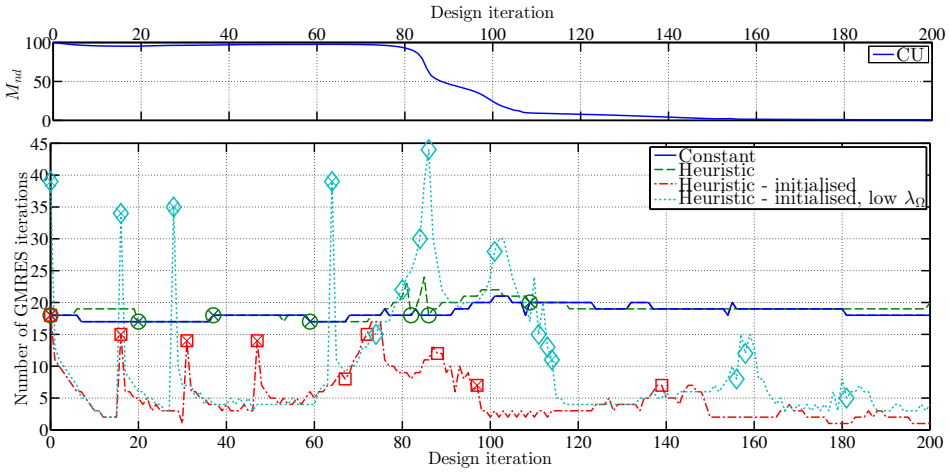

Figure 8: Development of the non-discreteness measure, M_{nd}, and a comparison of the number of GMRES iterations for the different basis computation strategies for a relative tolerance of $\varepsilon_{rel} = 10^{-5}$. The lines show the number of GMRES iterations for the various strategies and the symbols show when the basis is updated. Crosses highlight forced updates due to change in continuation parameter.

intermediate densities is relatively constant at the first and final stages of the optimisation process and that the contrast goes from low to high. The bottom plot of figure 8 shows the number of GMRES iterations during the optimisation process for the described strategies. It can be seen that the number of GMRES iterations is more or less constant throughout the optimisation process when the MsFEM basis is recomputed at every design iteration (denoted by "constant" in the legend). Keeping the development of M_{nd} in mind, this clearly supports the evidence that the MsFEM preconditioner is contrast-independent for a constant λ_Ω. This is of course an expensive approach and results in a total time for the entire design process of 768.5 seconds, which is significantly higher as compared to the reference time of 306.5 seconds when using the direct solver.

When adopting the heuristic update rule, the number of GMRES iterations can be seen to remain very close to constant throughout the optimisation process (denoted by "heuristic" in the legend). Here it can be seen that the MsFEM basis is only computed a total of 7 times during the 200 design iteration optimisation process, yielding a significantly lowered total

time for the entire design process of 444.0 seconds. It is interesting that of the 7 updates, the first is the initial basis computation and 4 are when the penalisation parameter, p, is updated which is a requirement made in the algorithm.

However, this is still significantly slower than the reference design process using the direct solver. Thus, as suggested, the displacement solution from the previous design iteration is used as the initial guess for the GMRES solver for the new design iteration. This significantly decreases the number of iterations needed to obtain the required accuracy when the design is slowly evolving. This is shown in figure 8 (denoted by "heuristic - initialised" in the legend), where it can be seen that 4-10 GMRES iterations are needed for large parts of the optimisation process, namely the initial and final stages, whereas significantly more are needed in the middle stage where the topology is forming and the design is changing rapidly.

By lowering the eigenvalue threshold, from $\lambda_\Omega = 6.69 \times 10^{-4}$ to $\lambda_\Omega = 4.64 \times 10^{-5}$, and thus decreasing the basis size, it is possible to decrease the computational work associated with the projection and a single GMRES iteration, as discussed in section 3.3. As can be seen in figure 8 (denoted by "heuristic - initialised, low λ_Ω" in legend), the required number of GMRES iterations generally increases as expected, however, despite this the resulting optimisation process only takes 298.7 seconds, which is 7.8 seconds faster than the reference time of 306.5 seconds when using the direct solver. This is an interesting observation and shows that the choice of the eigenvalue threshold, λ_Ω, is hardly trivial. Further investigation into determining the eigenvalue threshold using error estimates is the subject of future research.

All four final topologies are qualitatively very similar to the reference case and thus only a selected one is shown. The final topology for the optimisation process using the heuristic update rule combined with an initialised GMRES solver ("heuristic - initialised") is shown in figure 9c and can be seen to be very similar to the reference design obtained using the direct solver shown in figure 9a, except for hard corners as compared to rounded corners. The final compliance values are as follows; reference using direct solver: $C_{end} = 1.0967 \times 10^{-3}$, "constant": $C_{end} = 1.0812 \times 10^{-3}$, "heuristic": $C_{end} = 1.0762 \times 10^{-3}$, "heuristic - initialised": $C_{end} = 1.0649 \times 10^{-3}$, "heuristic - initialised, low λ_Ω": $C_{end} = 1.0787 \times 10^{-3}$. It is very interesting to note that all designs obtained using the MsFEM-GMRES approach actually performs better than the reference case, specifically 1.4%, 1.8%, 2.9% and 1.6%, respectively, for the four cases. However, this is coincidental due to the non-convexity of the

(a) Direct solver - reference case

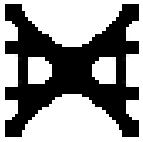

(b) Iterative solver - relaxed convergence criteria and seeded initial guess

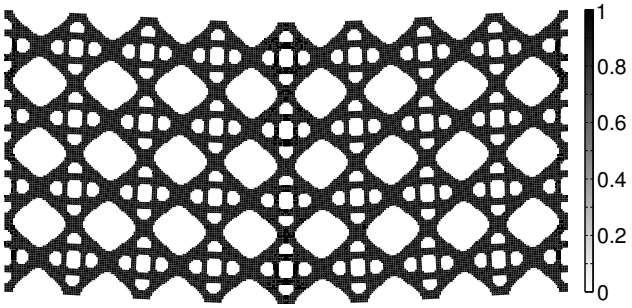

(c) Deformed macrostructure: Direct solver - reference case

Figure 9: Comparison of (a) the reference design by use of a direct solver and (b) the optimised design when using the proposed iterative method with relaxed convergence criteria. The deformed macrostructure is shown in (c) and is made up of 4 by 8 unit cells.

optimisation problem and can in no way be guaranteed in general.

5.3 Note on application to several thresholds for robust

Requiring robustness of the design response with respect to uniform erosion and dilation usually results in designs with similar topologies across the design realisations. This property can be utilised in the construction of the MsFEM preconditioner in order to further reduce the computational cost. Two strategies can easily be identified: 1) an offline building of the basis for the possible variations of the design and using it in an online optimisation process [44], and 2) building a basis for a representative design realisation and re-utilising it in the solutions of all other realisations. The first strategy is computationally expensive in the preparatory part (offline computations) compared to the second one, however it is expected to provide faster online computations. The second strategy does not require any additional modifications of the algorithm for building the coarse basis and provides a fast solution for the online optimisation process, as demonstrated later. Therefore, the second strategy is utilised in the examples presented in section 6.

The mass matrix in the eigenproblem defined, equation (9), acts as a filter which selects modes dominant within the solid region. As the dilated design realisation embeds all other design realisations, the basis for the dilated design will provide a good basis for all other design realisations. Therefore, the representative design is selected to be the dilated one. For small difference between the design realisations, which is the case here, the basis for the dilated design ensures a similar number of iterations for all solutions. If the difference between the design realisation is large, more rigorous selection criteria needs to be derived, which is left for future research.

6 Microstructural design examples

Throughout the microstructural design examples, the contrast in stiffness is constant and the Young's moduli of the solid and void are $E_{max} = 1$ and $E_{min} = 10^{-9}$, respectively.

6.1 Double clamped beam with concentrated load

The first numerical example is the double clamped beam with a concentrated load, shown in figure 1, as considered earlier in section 3.1. The problem is first investigated using the standard single design topology optimisation formulation and secondly using the robust formulation. The problem is optimised for an increasing number of unit cells in order to investigate the convergence of the compliance, as well as to compare the speed of the MsFEM-GMRES iterative solver, combined with the proposed basis re-utilisation scheme, to using the standard MATLAB direct solver. In order to make a fair comparison between the two, all of the timing runs have been executed on the same computer, where MATLAB has been restricted to use a single computational thread. The computer is a Dell Precision T7500 with two Intel Xeon X5650 CPUs and 96GB memory, running CentOS 6.4 and MATLAB 2013b.

6.1.1 Single design topology optimisation

The problem is first investigated using the standard single design topology optimisation formulation. The following parameters are used: $r_{min} = 4h$, $v_f = 0.5$, $\beta = 64$, $\eta = 0.5$, $\lambda_\Omega = 6.5 \times 10^{-4}$, $\varepsilon_{rel} = 10^{-6}$.

Figure 10 shows the optimised periodic microstructures for the double-clamped beam with a concentrated load in the centre. The microstructure is shown for increasingly smaller unit cell sizes. The number of unit cells across the height of the beam is characterised by $M_x = B/H$, which will be used throughout the paper. It is observed that the microstructure converges very fast and the overall topology does not qualitatively change for $M_x > 4$, therefore only a select few have been shown out of the investigated set, $M_x \in \{2, 4, 6, 8, 12, 16, 32, 64\}$. This trend was also observed in [16] for the same problem, although the optimised designs differ. This difference in topology is likely due to a difference in the effective length scale of the current approach and that of [16]; the former using density filtering combined with a projection at $\eta = 0.5$ and the latter using BESO.

For the optimised structures shown in figure 10, it is observed that the compliance increases as the number of unit cells is increased. This is attributed to the increased restriction of the design freedom imposed through the increased periodicity and is also observed in [16]. It is also seen that the MsFEM-GMRES approach gives qualitatively the same topologies as when

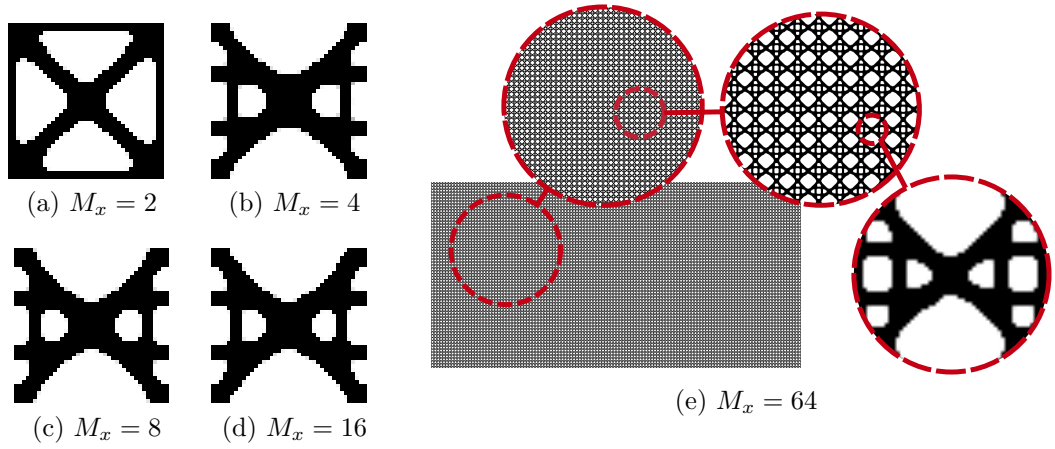

(a) $M_x = 2$ (b) $M_x = 4$ (c) $M_x = 8$ (d) $M_x = 16$ (e) $M_x = 64$

Figure 10: Optimised microstructures for the double-clamped beam with a concentrated load. The microstructure is shown for varying number of cells across the height, $M_x \in \{2, 4, 8, 16\}$, and the full macrostructure is shown for $M_x = 64$.

using a direct solver, however, slight deviations are observed in the compliance values.

From the numerical experiments of this section, it is also observed that in terms of average time per design iteration, the direct solver approach is faster for smaller problems, but that the MsFEM-GMRES approach becomes faster for larger problems, $M_x > 6$, as expected. For $M_x = 16$ the observed average times are 59.44 and 83.41 seconds for the MsFEM-GMRES and direct approaches, respectively. The relative difference is 28.7% and thus, the savings in time are quite substantial. Using a linear fit, the projected reduction in time for $M_x = 32$ and $M_x = 64$ is 41.0% and for 54.3%, respectively. However, for $M_x = 64$ it has been observed in practice that it takes approximately 26 hours for 200 design iterations. This yields a significantly lower average time per design iteration than the projected, yielding an expected reduction in time of 79.2%. Furthermore, a total time of 26 hours is quite impressive when considering that the optimisation was performed using MATLAB restricted to a single computational thread and taking into account that for $M_x = 64$, the total number of elements is 13,107,200 and the total number of degrees of freedom is 26,229,762. It should be noted that this problem was also solvable using a direct solver on the same machine, but the time taken for a single design iteration was so immense that the job was not allowed to

finish.

6.1.2 Robust topology optimisation

The problem is secondly investigated using the robust topology optimisation formulation. The following parameters are used: $r_{min} = 4h$, $v_f = 0.5$, $\beta_0 = 16$, $\beta_1 = 64$, $p_1 = 5$, $\eta = 0.3$, $\lambda_\Omega = 6.5 \times 10^{-4}$, $\varepsilon_{rel} = 10^{-6}$.

Figure 11 shows the optimised robust periodic microstructures for the double-clamped beam with a concentrated load in the centre. As for the single design case, it can be seen that the microstructure converges very fast and the overall topology does not qualitatively change for $M_x > 4$, therefore only a select few have been shown out of the investigated set, $M_x = \{2, 4, 6, 8, 12, 16\}$. It is interesting to note, that the robust topologies are very similar to the optimised topologies presented in [16]. The main difference is that the topologies presented in figure 11 have been optimised to be robust with respect to erosion and dilation of the design, which can be seen in figure 12 for $M_x = 16$. Comparing the non-robust and robust topologies, figures 10 and 11 respectively, it can be observed that the small horisontal and vertical bars in the left- and right-hand side of the unit cell of the non-robust design are not present in the robust designs. This is due to the length scale imposed by the robust formulation, as opposed to the lack of length scale for the non-robust result.

As for the non-robust case, it is observed that the compliance increases as the number of unit cells is increased. It is also seen that in terms of average time (total for all three realisations) per design iteration, the direct solver approach is faster for smaller problems, but the MsFEM-GMRES approach becomes faster for larger problems, $M_x > 8$. For $M_x = 16$ the observed

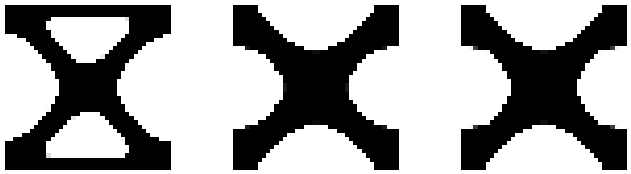

(a) $M_x = 2$ (b) $M_x = 4$ (c) $M_x = 16$

Figure 11: Optimised robust microstructures for the double-clamped beam with a concentrated load. The microstructures are shown for varying number of cells across the height, $M_x = \{2, 4, 16\}$.

(a) Eroded (b) Intermediate (c) Dilated

Figure 12: The three realisations of the optimised robust microstructures for the double-clamped beam with a concentrated load and $M_x = 16$.

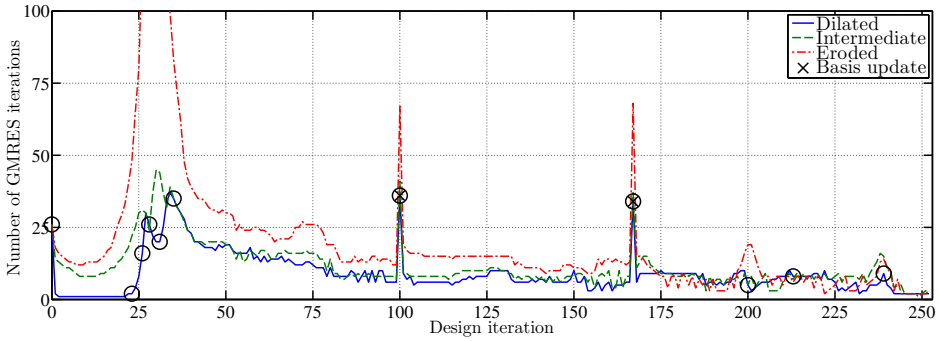

Figure 13: The number of GMRES iterations as a function of design iteration for the double clamped beam with a concentrated load using the robust formulation and $M_x = 16$.

average times were 188.91 and 246.68 seconds for the MsFEM-GMRES and direct approaches, respectively. The relative difference is 23.4% and thus, the savings in time are also quite substantial for the robust case. Using a linear fit, the projected reduction in time for $M_x = 32$ and $M_x = 64$ is 36.8% and 49.5%, respectively.

The observed cross-over point is higher for the robust case, which likely is due to a higher computational cost per design iteration. This higher computational cost is associated with the fact that the spectral coarse basis is formed for the dilated design and applied to all three realisations of the design. As expected, the iterative solution of the intermediate and eroded designs, using the spectral coarse basis based on the dilated design, needs more iterations to converge to the specified tolerance. For the adopted approach of using a moderately high β_0, the three realisations are significantly

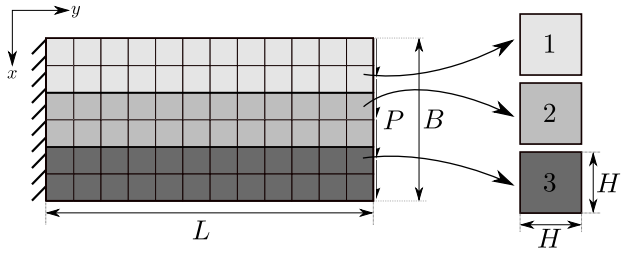

Figure 14: Illustration and dimensions of a cantilever beam subjected to a distributed force along the right-hand side. The beam is made up of three layers containing separate periodic microstructures with square unit cells.

different in the initial phase of the optimisation process when the topology is being formed. It is in this phase, where large differences in iteration numbers are observed, especially for the eroded case. Figure 13 shows the number of GMRES iterations during the optimisation process for $M_x = 16$. It can be seen that around design iterations 25-35, the number of GMRES iterations for the eroded design peaks at over 100. It is observed that the peak in GMRES iterations occurs at the same time as the designs are changing very quickly. When the topology has formed and the three realisations are similar, the performance can be seen to become much better and the difference in the number of GMRES iterations is relatively small. One could argue that it would be better to adopt the usual β-continuation, since these problems are not observed for this approach, because the realisations are very similar throughout the optimisation process. However, it has been observed that the presented approach needs significantly fewer design iterations to produce a competitive design and the approach is thus still faster, even with the increased computational time associated with the mismatch of the realisations. Improvement of the spectral basis for robust topology optimisation is a subject of future research.

6.2 Cantilever beam with distributed load

In order to demonstrate the methodology for design domains with multiple microstructural layers, a cantilever, as shown in figure 14, is investigated. The cantilever beam is subjected to a distributed force in the positive x-direction along the right-hand side. The height of the beam is set to $B = 1$, the length to $L = 2$ and the size of the force is set to $P = 0.01$ per unit

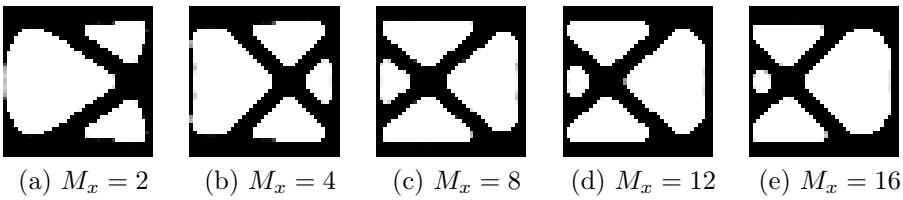

(a) $M_x = 2$ (b) $M_x = 4$ (c) $M_x = 8$ (d) $M_x = 12$ (e) $M_x = 16$

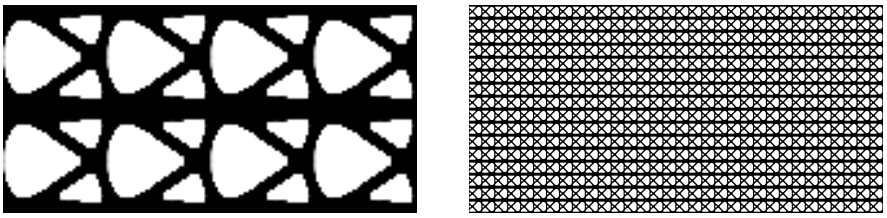

(f) $M_x = 2$ (g) $M_x = 16$

Figure 15: Optimised robust microstructures for the cantilever beam with a distributed load. The microstructures are shown for varying number of cells across the height, $M_x = \{2, 4, 8, 12, 16\}$. The full macrostructures are shown for $M_x = 2$ and $M_x = 16$.

length. The beam is made up of three layers with equal thickness, each subdivided into square coarse cells, with an edge length of H, containing a locally periodic microstructure. Each coarse cell is discretised with 40 by 40 linear finite elements. The following parameters are used: $r_{min} = 4h$, $v_f = 0.5$, $\beta_0 = 16$, $\beta_1 = 64$, $p_1 = 5$, $\lambda_\Omega = 6.5 \times 10^{-4}$, $\varepsilon_{rel} = 10^{-6}$.

The problem is first investigated for a single periodic microstructure and afterwards for a varying number of coarse cells across the thickness of the three layers. Finally, the problem is investigated for varying thickness of the outer two layers.

6.2.1 Single periodic robust microstructure

This problem was also investigated in [16], however, in the present work, the distributed load is in direct contact with the design domain, whereas a solid non-design region was included in [16] in order to allow for a direct comparison with results obtained through homogenisation. This is not done in the current work, in order to highlight the fact that by considering the full macroscopic structure, one actually obtains a design that automatically takes the boundary and loading conditions into account.

(a) Eroded (b) Intermediate (c) Dilated

Figure 16: The three realisations of the optimised robust microstructures for the cantilever beam with a distributed load and $M_x = 4$.

Figure 15 shows the optimised robust microstructures for the cantilever beam with a distributed load along the right-hand edge. It can be observed that, unlike for the double clamped beam problem of section 6.1, the microstructure does not converge within the observed range of unit cell refinement. It can be seen that the microstructure has the same overall topology from $M_x = 4 - 16$, but that the size of the holes are changing. By performing a cross-check analysis, this development in the topology appears to be due to local minima. This is not considered as important for the current study, as the aim of this paper is not to investigate the convergence of the unit cell design with respect to unit cell size .

Comparing the microstructures in figure 15 to those presented in [16], an important difference is the vertical bar at the right-hand side of the unit cell. This vertical bar is needed for the unit cells in contact with the distributed load and due to the imposed periodicity it exist in all unit cells within the beam, as seen in figures 15f and 15g. Of course, this leads to an increasingly sub-optimal design within the design domain, as the vertical bar is only needed at the very edge. But as already stated, it is seen as important to illustrate that the presented methodology, of taking the entire macrostructure into account, ensures a physically valid design with respect to boundary and loading conditions. The optimised microstructures can be seen to exhibit a combination of bending and shear stiffness, the first characterised by the horisontal bars and the latter characterised by the cross-like members. This is as expected, as forcing the microstructure to be the same all over the short beam, ensures that the microstructure needs to exhibit both characteristics.

Figure 16 shows the three realisations of the final optimised design for $M_x = 4$. Here it can be clearly seen, that the design is robust with respect to erosion and dilation. It is important to point out that the vertical bar

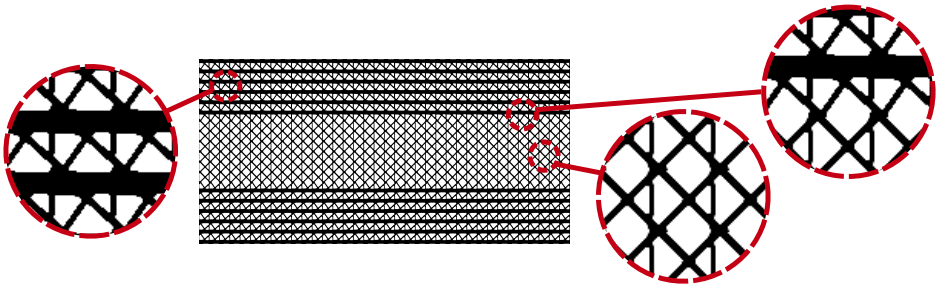

Figure 17: Optimised robust layered macrostructure for the cantilever beam with a distributed load for $M_x = 18$ and layer thickness $\left\{\frac{6}{18}, \frac{6}{18}, \frac{6}{18}\right\}$.

ensuring connection to the distributed load is ensured for the eroded design also, as seen in figure 16a.

6.2.2 Three-layered robust microstructure

One of the main benefits of analysing the full macrostructure with all microstructural details resolved, is the ability to design structures with varying microstructure and ensuring that these microstructures are connected. The presented implementation has the ability to optimise layered microstructures. The cantilever beam with a distributed load along the right-hand edge is split into three layers of equal thickness each made up of a number of periodic cells, as illustrated in figure 14.

Figure 17 shows the full macrostructure with optimised robust layered microstructures for $M_x = 18$ with layer thickness $\left\{\frac{6}{18}, \frac{6}{18}, \frac{6}{18}\right\}$, where the first entry denotes the relative thickness of the top layer, the second denotes the middle layer and the third denotes the bottom layer. It can clearly be seen that different microstructures exist in the three layers. As expected, the top and bottom layers exhibit microstructures with a high bending stiffness, characterised by the thick horisontal bars, whereas the middle layer exhibits a microstructure with a high shear stiffness, characterised by the cross structure. This is perfectly in line with the stress distribution of a short beam under bending; high axial stresses near the top and bottom and high shear stresses in the middle. Furthermore, due to the skew-symmetry of the stress distribution, the optimal topology should be symmetric [45] which is also observed for the optimised designs. It is important to note that no symmetry has been imposed during the optimisation process. It is observed that the average compliance does not change significantly with respect to increas-

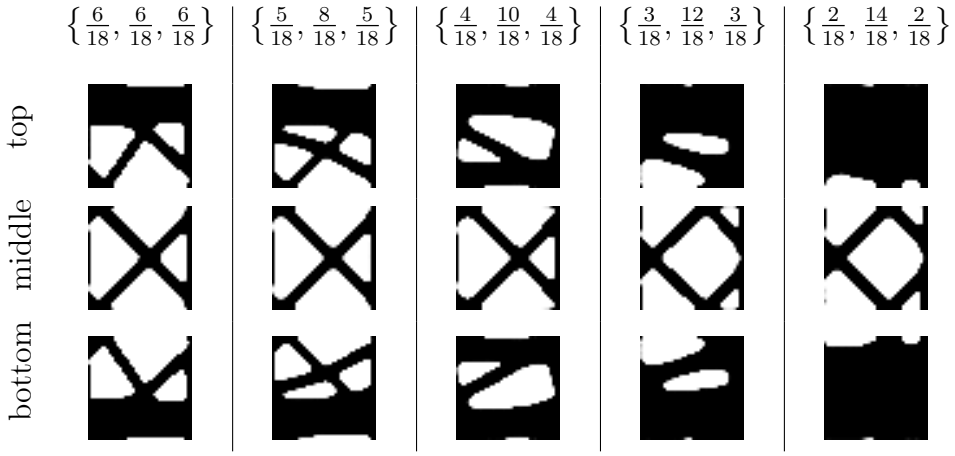

Figure 18: Optimised robust microstructures for the cantilever beam with a distributed load for three layers of varying outer layer thickness.

ing the number of coarse cells. Also, the microstructural topology does not change significantly and these are therefore not shown.

6.2.3 Three-layered robust microstructure of variable thickness

The three-layered cantilever beam with a distributed load along the right-hand edge is now investigated for varying thickness of the outer layers. Figure 18 shows the optimised robust microstructures for the three-layered cantilever beam with a distributed load for varying outer layer thickness. It can be seen that the microstructure of the middle layer does not change significantly when the thickness of the outer layers is reduced, in turn leading to an increase in the thickness of the middle layer. It should be noted, that although the topologies for $\left\{\frac{3}{18}, \frac{12}{18}, \frac{3}{18}\right\}$ and $\left\{\frac{2}{18}, \frac{14}{18}, \frac{2}{18}\right\}$ appear significantly different than for the others, they are merely the same topology as the others but shifted vertically half a coarse cell. However, the topology of the outer layers does change quite significantly as the thickness of the outer layers is reduced. It appears that the outer cells are becoming more dense and this is supported by looking at the local volume fraction of the outer coarse cells. Figure 19b clearly shows that as the thickness of the outer layers is reduced, the local volume fraction of material is increased in the outer layers. On the contrary, the local volume fraction of material in the middle layer is more or less constant. Figure 19a shows the average compliance and it can be seen that as

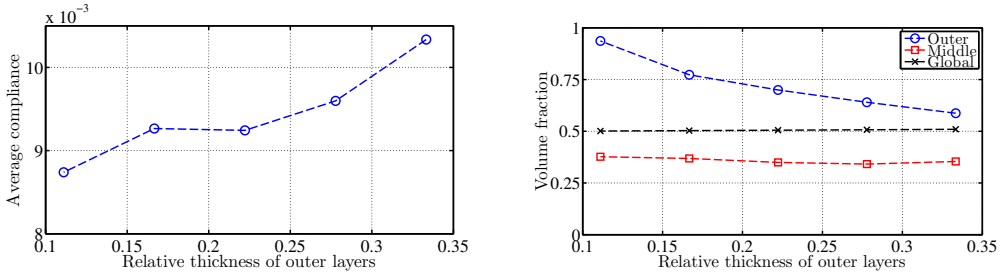

(a) Average compliance (b) Local volume fraction

Figure 19: Average compliance and local volume fractions as a function of the relative thickness of the outer layers for the three-layered cantilever beam with a distributed load.

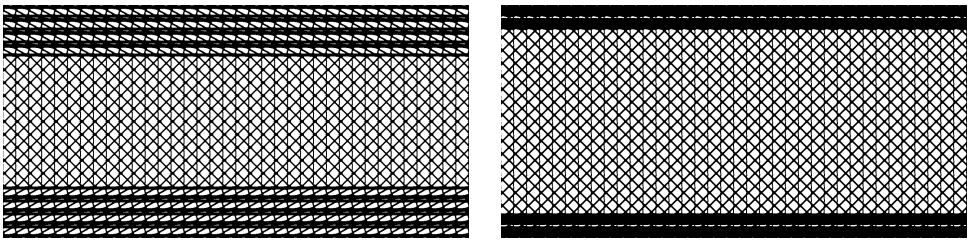

(a) $\left\{\frac{4}{18}, \frac{10}{18}, \frac{4}{18}\right\}$ (b) $\left\{\frac{2}{18}, \frac{14}{18}, \frac{2}{18}\right\}$

Figure 20: Optimised robust layered macrostructure for the cantilever beam with a distributed load for two selected layer thicknesses.

the thickness of the outer layers is reduced, the compliance decreases. Despite the compliance anomaly for $\left\{\frac{3}{18}, \frac{12}{18}, \frac{3}{18}\right\}$, likely due to a local minimum of the optimisation problem, it is concluded that it is beneficial to have thin outer layers with very high volume fractions and a thick middle layer with a relatively low volume fraction. This is perfectly in line with sandwich theory, where very thin and axially-stiff plates are joined to a low-density shear-stiff core in order to best resist the high axial stresses near the top and bottom and high shear stresses in the middle.

In order to fully appreciate the adaptivity of the approach to the varying outer layer thickness and the relation to sandwich theory, figure 20 shows the full macrostructure for two select layer thicknesses: $\left\{\frac{4}{18}, \frac{10}{18}, \frac{4}{18}\right\}$ and $\left\{\frac{2}{18}, \frac{14}{18}, \frac{2}{18}\right\}$. Here it can be clearly seen that the middle microstructure topology is essentially the same and that the outer layers become increasingly

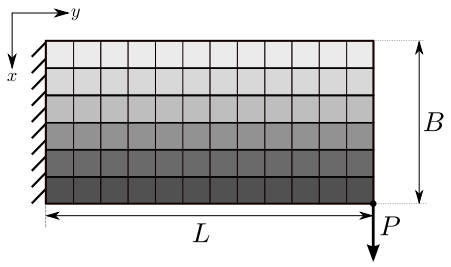

Figure 21: Illustration and dimensions of a cantilever beam subjected to a concentrated force at the bottom rightmost corner. Every layer has a unique microstructure which is periodic in the horisontal direction only.

dense.

6.3 Cantilever beam with concentrated load

In closing, in order to showcase the capabilities of the proposed methodology, a cantilever beam with a concentrated load at the lower right-hand corner is optimised where every layer of microstructure is different. Figure 21 shows the problem layout, where it can be seen that every layer has a unique microstructure which is periodic in the horisontal direction only. A small change is made to the filtering procedure, described in section 4.5, which is that internal periodicity in the vertical direction is no longer required. That is, the constraints applied on the horisontal edges of the microstructures, see figure 6b, are no longer imposed and the filtering domain essentially becomes a single vertical slice of the design domain with periodic boundary conditions in the horisontal direction. The following parameters are used: $r_{min} = 6h$, $v_f = 0.3$, $\beta_0 = 8$, $\beta_1 = 32$, $\eta \in \{0.3, 0.5, 0.7\}$, $\lambda_\Omega = 5 \times 10^{-4}$, $\varepsilon_{rel} = 10^{-6}$.

Figure 22 shows the optimised robust macrostructure for the cantilever beam with a concentrated load and twenty-four layers, $M_x = 24$. It can be seen that the microstructural details vary continuously throughout the thickness. Many layers are similar with small variations, due to gradual change of the principal strain directions. The top and bottom layers exhibits bending stiffness, where most of the inner layers exhibit shear stiffness. It can be noticed that there is a horisontal bar of material in a layer just below the middle. This is likely due to convergence to local minima, due to the rather aggressive continuation approach used. It has been observed that these features are less likely to appear when robustness is not required, however, they

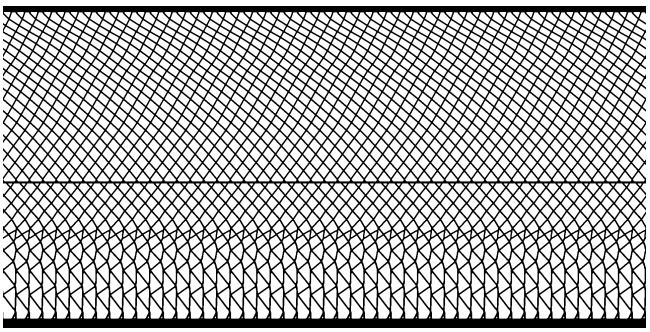

Figure 22: Optimised robust layered macrostructure for the cantilever beam with a concentrated load for $M_x = 24$.

do still appear at times and the problem increases with an increasing number of layers. The problem appears to be due to the propagation of the design information during the optimisation procedure, where the topologies of the local microstructural details are determined at very different speeds. Thus, the topologies of some layers are uniquely determined by the surrounding, already formed, layers.

7 Discussion

7.1 Need for manufacturable microstructures

The need for manufacturable microstructures of finite size is rather important. Current manufacturing processes, e.g. additive manufacturing, two photon polymerisation and photolithography, do not allow for the manufacture of microstructural details with length scales far below that of the specimen size. Furthermore, if the microstructural details are spatially-varying, it is crucial that connectivity is ensured throughout the manufactured structure in order to ensure the expected performance. It can be argued that going from a structure with disconnected, but optimal, microstructural details, as in [13] in the context of FMO, to modified connected microstructural details ensures a manufacturable and well-performing structure. But when

such modifications are made after the optimisation procedure itself, with no regard to the optimisation objective, the final structure may perform far from what was originally predicted. This may work to some extent for the forgiving compliance objective, but when moving to more advanced objectives or physics, the only rational approach is to include the requirement for connectivity in the original optimisation itself.

7.1.1 Connectivity constraints

It is important to note that small seemingly useless features can appear at the interfaces of the microstructural cells when dealing with multiple blocks. These small features are artefacts resulting from the very tough connectivity constraints imposed through the filtering procedure. In effect one imposes that all horisontal interfaces between any two given coarse cells should be the same. Generally it is observed that the microstructures that develop a clear topology first, that is the ones with the largest sensitivities, have a dominant influence on the interfaces. This is as expected, but the filtering procedure ensures connected microstructures that have as smoothly connected interfaces as possible in order to ensure manufacturability. When the design is further relaxed into multiple layers with unique microstructures, these artefacts do not occur because internal periodicity is not required in the vertical direction.

7.2 Microstructure restriction

In this article, the microstructural details have been forced to be periodic or in layers. The ideal situation is, of course, the ability of the microstructural details to change in all directions. Forcing the design to have microstructural details everywhere, essentially imposes a maximum length scale on the resulting members. Providing the possibility of unrestricted microstructural details is trivial for the presented approach, as the optimisation problem reverts back to the original unrestricted topology optimisation problems. In this case, all length scales above the minimum length scale imposed by the robust approach are allowed in the design space and thus a significantly better performing structure is to be expected than when restricting the design space using forced microstructural details. If one for instance wants microstructural details everywhere, one can imposed a maximum length scale on the design using local lower-bound constraints on the void volume fraction, similar to

the approach presented in [17].

7.3 Spectral MsFEM as a solver

Most multiscale finite element methods have been developed for use as an approximate solver, rather than as a preconditioner. It is also possible to use the presented spectral MsFEM as a solver for high-contrast linear elasticity problems. As the error convergence studies presented in section 3.1 suggest, a lot of eigenmodes are needed to approximate the true solution well. Using only a few modes yields results similar to homogenisation and thus designs can end up being massively disconnected, as discussed in the original work [18]. By increasing the number of modes slightly, one can increase the sensitivity of the structural response to disconnected members. However, for microstructural design problems, these low-dimension bases require eigenfunction derivatives for the optimisation to proceed correctly, which is computationally expensive. Hence, using the presented spectral MsFEM as a solver is left as a subject for future research.

7.4 Future work

For the presented problems, the addition of local constraints on the minimum void volume fraction in each unique microstructure cell is a cheap way to ensure a maximum length scale, similar to [17]. Of course, the number of local constraints increases with the amount of unique microstructures, that is the number of layers. This approach has been tested successfully, however, it does not show significant differences for the presented problems with relatively low volume fractions. The extension to other local constraints will be investigated in future work.

A full comparison of the presented approach to homogenisation for varying microstructural details is part of current research, which will include investigations on the possibility of using the presented methodology to find the optimal design for cores of sandwich beam, plate and shell structures. The extension of the MsFEM-GMRES solver to three-dimensional general topology optimisation problems in a large scale parallel framework [27, 46] is currently being pursued. Here the computational performance of the method is expected to excel due to the inherent parallisation properties of the decoupled local problems.

8 Conclusion

The presented approach treats the optimisation of microstructural details with respect to the macroscopic response, while resolving all microstructural details as compared to the homogenisation approach. This results in an optimisation that takes localised behaviour, such as boundary conditions and forces, into account and ensures microstructures tailored for a specific application rather than specific properties. The approach is fully flexible and can treat microstructural details with and without clear separation of length scale compared to macrostructure. Finally, the approach yields robust and manufacturable solutions, with smoothly connected and continuously varying microstructural details in one direction. A further advantage of the proposed methodology, is that the application of the MsFEM-GMRES solver allows for huge problems to be solved in MATLAB and that it exhibits contrast-independent behaviour ensuring fast computations for topology optimisation problems.

9 Acknowledgements

The authors would like to thank Erik Andreassen for providing MATLAB subroutines for homogenisation and the TopOpt group for useful comments and discussions during the preparation of the work.

Both authors were funded by Villum Fonden through the NextTop project, as well as the EU FP7-MC-IAPP programme LaScISO.

References

[1] O. Sigmund, Materials with prescribed constitutive parameters: An inverse homogenization problem, International Journal of Solids and Structures 31 (17) (1994) 2313 – 2329. doi:10.1016/0020-7683(94)90154-6.

[2] O. Sigmund, Tailoring materials with prescribed elastic properties, Mechanics of Materials 20 (4) (1995) 351 – 368. doi:10.1016/0167-6636(94)00069-7.

[3] O. Sigmund, S. Torquato, Composites with extremal thermal expansion coefficients, Applied Physics Letters 69:21 (1996) 3203–3205.

[4] C. S. Andreasen, O. Sigmund, Multiscale modeling and topology optimization of poroelastic actuators, Smart Materials and Structures 21 (6) (2012) 065005. doi:10.1088/0964-1726/21/6/065005.

[5] M. P. Bendse, O. Sigmund, Topology Optimization: Theory, Methods and Applications, Springer, 2003, iSBN: 3-540-42992-1.

[6] M. P. Bendsøe, N. Kikuchi, Generating optimal topologies in structural design using a homogenization method, Computer Methods in Applied Mechanics and Engineering 71 (2) (1988) 197–224. doi:10.1016/0045-7825(88)90086-2.

[7] M. P. Bendsøe, Optimal shape design as a material distribution problem, Structural optimization 1 (4) (1989) 193–202. doi:10.1007/BF01650949.

[8] H. Rodrigues, J. Guedes, M. Bendsoe, Hierarchical optimization of material and structure, Structural and Multidisciplinary Optimization 24 (1) (2002) 1–10. doi:10.1007/s00158-002-0209-z.

[9] P. G. Coelho, P. R. Fernandes, H. C. Rodrigues, J. B. Cardoso, J. M. Guedes, Numerical modeling of bone tissue adaptation - a hierarchical approach for bone apparent density and trabecular structure, Journal of Biomechanics 42 (2009) 830–837.

[10] P. G. Coelho, J. B. Cardoso, P. R. Fernandes, H. C. Rodrigues, Parallel computing techniques applied to the simultaneous design of structure and material, Advances in Engineering Software 42 (2011) 219–227.

[11] P. B. Nakshatrala, D. A. Tortorelli, K. B. Nakshatrala, Nonlinear structural design using multiscale topology optimization. Part I: Static formulation, Computer Methods in Applied Mechanics and Engineering 261-263 (2013) 167–176.

[12] L. Xia, P. Breitkopf, Concurrent topology optimization design of material and structure within nonlinear multiscale analysis framework, Computer Methods in Applied Mechanics and Engineering 278 (2014) 524–542. doi:http://dx.doi.org/10.1016/j.cma.2014.05.022.

[13] F. Schury, M. Stingl, F. Wein, Efficient two-scale optimization of manufacturable graded structures, SIAM Journal on Scientific Computing 34 (6) (2012) B711–B733.

[14] A. Radman, X. Huang, Y. Xie, Topology optimization of functionally graded cellular materials, Journal of Materials Science 48 (4) (2013) 1503–1510.

[15] Y. Xie, Z. Zuo, X. Huang, J. Rong, Convergence of topological patterns of optimal periodic structures under multiple scales, Structural and Multidisciplinary Optimization 46 (1) (2012) 41–50.

[16] Z. H. Zuo, X. Huang, X. Yang, J. H. Rong, Y. M. Xie, Comparing optimal material microstructures with optimal periodic structures, Computational Materials Science 69 (2013) 137–147.

[17] J. Guest, Imposing maximum length scale in topology optimization, Structural and Multidisciplinary Optimization 37 (2009) 463–473, 10.1007/s00158-008-0250-7.

[18] B. S. Lazarov, Topology optimization using multiscale finite element method for high-contrast media, in: I. Lirkov, S. Margenov, J. Wasniewski (Eds.), Large-Scale Scientific Computing, Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2014, pp. 339–346. doi:10.1007/978-3-662-43880-0_38.

[19] I. Babuška, G. Caloz, J. E. Osborn, Special finite element methods for a class of second order elliptic problems with rough coefficients, SIAM Journal on Numerical Analysis 31 (4) (1994) pp. 945–981.

[20] I. Babuška, J. E. Osborn, Generalized finite element methods: Their performance and their relation to mixed methods, SIAM Journal on Numerical Analysis 20 (3) (1983) pp. 510–536.

[21] T. Hou, X. Wu, A multiscale finite element method for elliptic problems in composite materials and porous media, Journal Of Computational Physics 134 (1) (1997) 169–189.

[22] Y. Efendiev, T. Y. Hou, Multiscale Finite Element Methods: Theory and Applications, Springer, 2009.

[23] M. Buck, O. Iliev, H. Andrä, Multiscale finite element coarse spaces for the application to linear elasticity, Central European Journal of Mathematics 11 (4) (2013) 680–701. doi:10.2478/s11533-012-0166-8.

[24] J. Galvis, Y. Efendiev, Domain decomposition preconditioners for multiscale flows in high-contrast media, Multiscale Modeling & Simulation 8 (4) (2010) 1461–1483. doi:10.1137/090751190.
URL http://link.aip.org/link/?MMS/8/1461/1

[25] Y. Efendiev, J. Galvis, X.-H. Wu, Multiscale finite element methods for high-contrast problems using local spectral basis functions, Journal of Computational Physics 230 (4) (2011) 937 – 955. doi:10.1016/j.jcp.2010.09.026.

[26] D. Braess, Finite Elements: Theory, Fast Solvers, and Applications in Solid Mechanics, Cambridge University Press, 2007.

[27] N. Aage, B. Lazarov, Parallel framework for topology optimization using the method of moving asymptotes, Structural and Multidisciplinary Optimization 47 (4) (2013) 493–505. doi:10.1007/s00158-012-0869-2.

[28] O. Amir, N. Aage, B. Lazarov, On multigrid-CG for efficient topology optimization, Structural and Multidisciplinary Optimization 49 (5) (2014) 815–829. doi:10.1007/s00158-013-1015-5.
URL http://dx.doi.org/10.1007/s00158-013-1015-5

[29] P. S. Vassilevski, Multilevel Block Factorization Preconditioners: Matrix-based Analysis and Algorithms for Solving Finite Element Equations, Springer, New York, 2008.

[30] Y. Efendiev, J. Galvis, P. Vassilevski, Spectral element agglomerate algebraic multigrid methods for elliptic problems with high-contrast coefficients, in: Y. Huang, R. Kornhuber, O. Widlund, J. Xu (Eds.), Domain Decomposition Methods in Science and Engineering XIX, Vol. 78 of Lecture Notes in Computational Science and Engineering, Springer Berlin Heidelberg, 2011, pp. 407–414.

[31] Y. Efendiev, J. Galvis, R. Lazarov, J. Willems, Robust domain decomposition preconditioners for abstract symmetric positive definite bilinear forms, ESAIM: Mathematical Modelling and Numerical Analysis 46 (2012) 1175–1199.

[32] Y. Saad, M. Schultz, GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems, SIAM Journal on Scientific and Statistical Computing 7 (3) (1986) 856–869. doi:10.1137/0907058.

[33] M. Brezina, P. S. Vassilevski, Smoothed aggregation spectral element agglomeration AMG: SA-ρAMGe, in: I. Lirkov, S. Margenov, J. Waniewski (Eds.), Large-Scale Scientific Computing, Vol. 7116 of Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2012, pp. 3–15. doi:10.1007/978-3-642-29843-1_1.

[34] K. Svanberg, The method of moving asymptotes - a new method for structural optimization, International Journal for Numerical Methods in Engineering 24 (2) (1987) 359–373. doi:10.1002/nme.1620240207.

[35] P. Michaleris, D. A. Tortorelli, C. A. Vidal, Tangent operators and design sensitivity formulations for transient non-linear coupled problems with applications to elastoplasticity, International Journal for Numerical Methods in Engineering 37 (14) (1994) 2471–2499. doi:10.1002/nme.1620371408.

[36] J. K. Guest, J. H. Prvost, T. Belytschko, Achieving minimum length scale in topology optimization using nodal design variables and projection functions, International Journal for Numerical Methods in Engineering 61 (2) (2004) 238–254. doi:10.1002/nme.1064.

[37] F. Wang, B. S. Lazarov, O. Sigmund, On projection methods, convergence and robust formulations in topology optimization, Structural Multidisciplinary Optimization 43 (6) (2011) 767–784. doi:10.1007/s00158-010-0602-y.

[38] B. Bourdin, Filters in topology optimization, International Journal for Numerical Methods in Engineering 50 (9) (2001) 2143–2158. doi:10.1002/nme.116.

[39] O. Sigmund, Morphology-based black and white filters for topology optimization, Structural Multidisciplinary Optimization 33 (4-5) (2007) 401–424. doi:10.1007/s00158-006-0087-x.

[40] B. S. Lazarov, O. Sigmund, Filters in topology optimization based on helmholtz-type differential equations, International Journal for Numerical Methods in Engineering 86 (2011) 765–781.

[41] B. S. Lazarov, M. Schevenels, O. Sigmund, Topology optimization considering material and geometric uncertainties using stochastic collocation methods, Structural and Multidisciplinary Optimization 46 (2012) 597–612.

[42] J. K. Guest, A. Asadpoure, S.-H. Ha, Eliminating beta-continuation from heaviside projection and density filter algorithms, Structural and Multidisciplinary Optimization 44 (4) (2011) 443–453. doi:10.1007/s00158-011-0676-1.

[43] O. Amir, O. Sigmund, B. S. Lazarov, M. Schevenels, Efficient reanalysis techniques for robust topology optimization, Computer Methods in Applied Mechanics and Engineering 245/246 (2012) 217 – 231.

[44] Y. Efendiev, J. Galvis, T. Y. Hou, Generalized multiscale finite element methods (gmsfem), Journal of Computational Physics 251 (0) (2013) 116 – 135. doi:http://dx.doi.org/10.1016/j.jcp.2013.04.045.

[45] G. I. Rozvany, On symmetry and non-uniqueness in exact topology optimization, Structural and Multidisciplinary Optimization 43 (3) (2011) 297–317. doi:10.1007/s00158-010-0564-0.

[46] N. Aage, E. Andreassen, B. S. Lazarov, Topology optimization using PETSc: An easy-to-use, fully parallel, open source topology optimization framework, Structural and Multidisciplinary Optimization (2014) available online.